HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P. O. Box 272400
Fort Collins, Colorado 80527-2400

IFW AF

PATENT APPLICATION

ATTORNEY DOCKET NO. 10016693-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Donald C. Soltis, Jr.

Confirmation No.: 8480

Application No.: 10/092,645

Examiner: Pierre M. Vital

Filing Date: March 6, 2002

Group Art Unit: 2188

Title: System And Method For Dynamic Processor Core And Cache Partitioning On Large-Scale Multithreaded, Multiprocessor Integrated Circuits

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Sir:

Transmitted herewith in **triplicate** is the Appeal Brief in this application with respect to the Notice of Appeal filed on **June 23, 2004**.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) $330.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

( ) (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d) for the total number of months checked below:

( )	one month	$110.00
( )	two months	$420.00
( )	three months	$950.00
( )	four months	$1480.00

( ) The extension fee has already been filled in this application.

(X) (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account **08-2025** the sum of **$330.00**. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

(X) I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, Alexandria, VA 22313-1450. Date of Deposit: **Aug. 17, 2004**

OR

( ) I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number_____ on _____

Number of pages: **35**

Typed Name: **Chalynda Renz**

Signature: Chalynda Renz

Respectfully submitted,

Donald C. Soltis, Jr.

By _____

Peter C. Knops

Attorney/Agent for Applicant(s)
Reg. No. **37,659**

Date: **Aug. 17, 2004**

Telephone No.: **(816) 460-5826**

PATENT
Attorney Docket No. 10016693-1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Appellant:	Soltis
Serial No.:	10 / 092,645
Filed:	03/06/2002
Group Art Unit:	2188
Examiner:	Pierre M. Vital
For:	System And Method For Dynamic Processor Core And Cache Partitioning On Large-Scale Multithreaded, Multiprocessor Integrated Circuits

BRIEF

PURSUANT TO 37 C.F.R. § 1.192

I. Real Party in Interest

The real party in interest for this appeal is Hewlett-Packard Development Company, L.P. (HPDC), a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, L.L.C. Evidence of this assignment, which was recorded on June 18, 2003, may be found at reel/frame 013776/0928.

II. Related Appeals and Interferences

No other appeals or interferences are currently known to Appellant that will directly affect, be directly affected by, or have a bearing on the decision to be rendered by the Board of Patent Appeals and Interferences in the present appeal.

III. Status of Claims

Claims 1–11 stand rejected, and are pending for consideration in this appeal.

IV. Status of Amendments

No amendments to the claims were filed subsequent to the Final Office Action issued on 03/23/2004.

V. Summary of the Invention

The following summary is an edited excerpt based on the Summary of the Invention section in Appellant's specification, paragraph numbers 16, 17, and 18:

A high performance multiple-processor integrated circuit has multiple cache units and multiple instruction fetch and decode units, where each instruction fetch and decode unit is associated with a real or virtual processor. The integrated circuit also has at least one dynamically allocable cache unit, and may have additional cache units directly connected to particular processors.

The integrated circuit also has high-speed interconnect that permits use of the dynamically allocable cache units by more than one real or virtual processor, and a cache allocation register. Fields of a cache allocation register determine which real or virtual processor is permitted to access each dynamically allocable cache unit.

In an embodiment, the dynamically allocable cache units form part of the second level of cache in the system.

VI. Issues

1. *Whether the combination of* Otterness et al, U.S. Patent 6,460,122, Akashi et al, U.S. Patent 6,438,653, and of Liao et al, U.S. Patent 6,681,296 *fails to teach or suggest every element recited in Appellant's independent claims 1, and 5 in the context of main memory cache on a multiprocessor integrated circuit, and therefore fails to render the claims obvious under 35 U.S.C. sec. 103.*

2. *Whether* Otterness *can be relied upon by the Examiner for the element of a cache block allocable to cache controllers when the only allocable element disclosed in* Otterness *is a cache line.*

VII. Grouping of Claims

Group I: Claims 1–11 stand together.

VIII. Argument

A. REASONS FOR ALLOWABILITY OF THE CLAIMS

1. The combination of the Otterness, Akashi, and Liao references fails to teach or suggest every element recited in Appellant's independent claims 1 and 5 in the context of main memory cache, and therefore fails to render the claimed invention obvious under 35 U.S.C. 103.

Appellant's claimed invention, as recited in claim 1, is a multiprocessor integrated circuit having at least two processors, each having an associated cache controller for caching references to a random access memory (RAM). There are multiple cache memory blocks on the processor integrated circuit, some of which are allocable for use by selected cache controllers.

Appellant's claim 5 recites a related method for allocating cache blocks on the multiple-processor integrated circuit of claim 1.

For a claim to be obvious under 35 U.S.C. 103 in view of cited art, it is necessary to find each of the elements, or strong suggestions of the elements, of the claim in the art, together with a suggestion in the art that these elements be combined to produce the claimed invention.

Otterness describes a disk controller subsystem, the controller having multiple embedded processors, for connecting multiple processor-with-memory systems to a redundant array of disk drives. Otterness does not describe a multiprocessor integrated circuit having a memory cache, as claimed in Appellant's independent claim 1, and does not describe a method

for allocating cache on a multiprocessor integrated circuit as recited in Appellant's independent claim 5. The disk controller in Otterness incorporates cache memory for caching "RAID stripe cache line" data read from the array of disks attached to the disk controller. Intrinsic to operation of Otterness is allocation of small portions of cache memory to "RAID stripe cache lines" as they are read from the disk array.

In Otterness, cache memory is allocated to "RAID stripe cache line" data after data requested by a first processor has been found to be absent in cache. Since the cache memory is allocated to data, not to processors, a second processor requesting the same data will receive the data from the same cache memory location. In Appellant's cache, cache memory is allocated to processors in advance of the processor requesting data.

While some RAID stripe cache lines may be requested by a first processor-with-memory system, and others may be requested by a second processor-with-memory system, Otterness lacks the element or concept of repartitioning a system (independent claim 5), or a resource allocation controller (independent claim 1) to allocate memory blocks capable of holding multiple cache lines to *processors* or *caches* as recited in Appellant's independent claims.

As observed in Appellant's response to the initial office action, Appellant notes that US patent 6,438,658 is to Baglia; appellant believes that whenever the Examiner uses this number while naming Akashi, patent no. 6,438,653 to Akashi is intended.

Akashi describes a main memory cache having a cache-coherence maintenance capability. Akashi provides a mechanism for allocating cache lines to fetched data. Some cache lines may correspond to data requested by a first processor-with-memory system, and others by a second processor-with-memory system, but Akashi also lacks the element or concept of

allocating reallocable memory blocks to *processors* or *caches* as recited in Appellant's independent claims 1 and 5.

Locating cache lines having current data in other caches as per Akashi is a fundamentally different process from the claimed process of allocating memory blocks to caches.

Liao describes a cache providing for two cache blocks, where a first block may be "locked," such that cache lines stored in a first block are not replaced by recently-fetched data, and lines stored in a second block are replaced by recently-fetched data. A configuration register is provided for controlling whether the first block operates normally, or is "locked."

Of the cited art, only Liao provides a cache memory block larger than a cache line, and only Liao provides a cache configuration register for controlling cache memory blocks.

Liao however fails to provide the element or concept of allocating memory blocks capable of holding multiple cache lines to *processors* or *caches* as recited in Appellant's independent claims 1 and 5.

All three patents of cited art thus fail to provide the key concept of allocating memory blocks capable of holding multiple cache lines to *processors* or *caches* as required by the claims; thus key elements of the claims are not present in the cited art.

Furthermore, even if memory blocks capable of holding multiple cache lines were actually allocated to *processors* or *caches* in the cited art, there is no suggestion that these memory blocks be combined into a second level cache on a repartitionable processor integrated circuit as recited in Appellant's claims 1 and 5 to produce the claimed invention.

Therefore, for at least the above reasons, the combination of the Otterness, Akashi, and Liao references is not sufficient to show obviousness under 35 U.S.C. 103, since the references fail to teach or suggest every element of Appellant's

invention, as recited in each one of the independent claims 1 and 5, in the context of a processor integrated circuit.

2. *Otterness* can not be relied upon by the Examiner for the element of a cache block allocable to cache controllers when the only allocable element disclosed in *Otterness* is a cache line.

Otterness features multiple host processors and multiple controller processors.

In Otterness, a small unit, or cache line, of memory is allocated to each "RAID stripe cache line" after data requested by a first host processor has been found absent in cache. Data is then fetched from the disk systems into the cache line. Since the cache memory is allocated to data, *not to host processors*, a second processor requesting the same data will receive the data from the same cache memory location.

In paragraph 33 of Appellant's application, it is stated that allocation of cache to processors is performed through a partition allocation table. In paragraph 29, reference is made to reallocable cache blocks that incorporate tag memory and interconnect interfaces. In paragraph 22 of Appellant's application, reference is made to clearing multiple cache lines in a block when a cache block is reassigned to a different processor. These references are only consistent with an interpretation of the term "cache memory block" used in Appellant's independent claims 1 and 5 as a unit of cache memory larger than the cache lines of Otterness, Akashi, and Liao.

Therefore the combination of Otterness, Akashi, and Liao is deficient with respect to showing each element of Appellant's claims 1 and 5, since these references separately or together fail to teach or suggest the element of a cache block allocable to cache controllers. Accordingly, the combination of Otterness, Akashi, and Liao cannot show obviousness of Appellant's claimed invention, as recited in claims 1 and 5.

B. Grouping of Claims

The claims in Group I stand or fall together.

IX. Conclusion

The Otterness, Akashi, and Liao references form basis for the rejection of each of Appellant's claims in the final action of record. Because Appellant has shown that these references are deficient with respect to teaching or suggesting at least one element (e.g., cache blocks allocable to cache controllers) in each of Appellant's independent claims, every rejection of these claims must therefore fail, whether based on 35 U.S.C. sec.102 *or* sec. 103, in view of the fact that this element, at least, is not taught or suggested by any of the cited references.

Claims in dependent form shall be construed to include all the limitations of the claim incorporated by reference into the dependent claim. 37 CFR 1.175. Therefore, dependent claims 2-4 and 6-11 are also allowable, since each of these claims depends from and incorporates all the limitations of an allowable independent claim.

Appellant therefore believes that all pending claims are allowable over the cited art because, *inter alia*, the present Office Action fails to establish a *prima facie* case of unpatentability for any of the claims. Such a *prima facie* case is non-existent because, among other things, (1) there is no prior art presented that individually or cumulatively teaches or suggests each of the elements of any of Appellant's claims, and (2) no proper motivation is provided to modify the teachings of the Steely reference. During patent examination the PTO bears the initial burden of presenting a *prima facie* case of unpatentability. *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). If the PTO fails to meet this burden, then the appellant is entitled to the patent. *In re Glaug*, 62 USPQ2d 1151 (Fed. Cir. 2002).

Appellant believes that claims 1–11 fully comply with 35 U.S.C. § 112 and, for at least the above reasons, are not obvious in view of the cited art to one skilled in the art having knowledge thereof. Accordingly, Appellant respectfully submits that claims 1–11 are patentable over the prior art and respectfully requests this Board to so indicate.

Respectfully submitted,

Peter C. Knops, Reg. No. 37,659
LATHROP & GAGE LC
2345 Grand Boulevard, Suite 2400
Kansas City, MO 64108
(816) 460-5826 Telephone
(816) 292-2001 Facsimile

X. APPENDIX OF CLAIMS ON APPEAL

What is claimed is:

1. (Previously Amended) A processor integrated circuit capable of executing more than one instruction stream comprising:

a first processor, coupled to fetch instructions and access data through a first cache controller;

a second processor, coupled to fetch instructions and access data through a second cache controller;

a plurality of cache memory blocks, each containing memory;

a high-speed interconnect coupling the plurality of cache memory blocks to the first and second cache controllers such that at least one allocable cache memory block is capable of being used by the first and second cache controllers; and

a resource allocation controller coupled to determine an accessing cache memory controller selected from the group consisting of the first and second cache memory controllers, whereby the accessing cache memory controller is allowed to access the allocable cache memory block,

wherein the cache memory blocks are usable by the cache controllers to store data and instructions fetched from a random-access memory.

2. (Original) The processor integrated circuit of Claim 1, further comprising a plurality of first level cache systems, wherein the first processor fetches instructions and accesses data from the first cache controller through a first first level cache system, and wherein the second processor fetches instructions and accesses data from the second cache controller through a second first level cache system.

3. (Original) The processor integrated circuit of Claim 1, wherein the cache memory blocks further comprise cache tag memory.

4. (Original) The processor integrated circuit of Claim 1, wherein each cache controller is provided with cache hit rate monitoring apparatus.

5. (Previously Amended) A method of dynamically allocating cache on a multiple-processor integrated circuit, where the multiple processor integrated circuit is used in a partitionable multiple-processor system and comprises:

a plurality of processors each coupled to receive instructions from a first level cache associated therewith,

a plurality of allocable upper level cache memory blocks,

interconnect apparatus for transmitting cache misses at each first level cache to upper level cache memory blocks assigned thereto, and

allocation apparatus for assigning upper level cache memory blocks to processors;

the method comprising the steps of:

monitoring past cache performance associated with processors and partitions;

determining desired processor to partition and upper level cache block allocations to processors; and

repartitioning the system, the step of repartitioning the system including allocation of upper level cache blocks to processors of at least one of the multiple processor integrated circuits~~.~~,

wherein the upper level cache blocks are usable to store data and instructions fetched from a random-access main memory.

6. (Original) The method of claim 5, wherein the upper level cache blocks are second level cache blocks.

7. (Original) The method of Claim 5, further comprising the step of billing customers according to processor time and allocated cache.

8. (Original) The method of Claim 5, wherein the multiple processor integrated circuit further comprises a plurality of non-allocable cache memory blocks.

9. (Original) The method of Claim 5, wherein the interconnect apparatus further comprises a plurality of upper level cache controllers, and where each upper level cache controller is capable of controlling operation of the allocable cache memory blocks as a writeback cache.

10. (Original) The method of Claim 5, wherein each upper level allocable cache block further comprises tag memory and cache data memory.

11. (Previously Amended) The method of Claim 10, further comprising the steps of stopping execution of operating systems in each partition, and restarting execution of operating systems in each partition, and wherein the system is capable of being repartitioned without rebooting each operating system.

9. (Original) The method of Claim 5, wherein the interconnect apparatus further comprises a plurality of upper level cache controllers, and where each upper level cache controller is capable of controlling operation of the allocable cache memory blocks as a writeback cache.

10. (Original) The method of Claim 5, wherein each upper level allocable cache block further comprises tag memory and cache data memory.

11. (Previously Amended) The method of Claim 10, further comprising the steps of stopping execution of operating systems in each partition, and restarting execution of operating systems in each partition, and wherein the system is capable of being repartitioned without rebooting each operating system.